

Systems Software Services Technical Bulletin

Number: 0015
Issued Date: April 26, 1982
Effective Date:
Section/Groups:
Submitted By:
Approved By:

VSAM File Define Time Considerations

If you are working with VSAM files, or anticipate doing so, the following generalized guidelines may be of help in setting up and 'tuning' your files for better performance. Again, emphasis is placed on generalized guidelines as each application will be different in needs and options.

SPEED

In 98% of the cases, specify SPEED—don't default to RECOVERY. This option applies only to the loading of the file.

NOWRITECK

(Is the default)—Don't override! Incurs large amounts of overhead.

SPACE ALLOCATION

Specify CYLINDERS, not TRACKS. Specify enough primary space to contain the entire dataset. Specify approximately 20% of primary for secondary.

FREESPACE

If inserts are going to be evenly distributed across the file, increase CI FREESPACE (FSPC) % and decrease CA FSPC %. If inserts will be clustered, decrease CI FSPC % and increase CA FSPC %. Caution: If clustered inserts will be at the end of an ascending sequential file, you will avoid many CI splits and performance degradation by first adding a record with a key of high-values (either hex "FF..." or numeric 9's).

IMBED

If you anticipate having a large number of retrievals, consider using EMBED, which can improve performance. Be aware, however, that using IMBED will cause a slightly higher number of CI splits.

UNIQUE

Specify UNIQUE rather than default to SUBALLOCATION. New extensions of VSAM software will soon require the use of UNIQUE.

NAMING CONVENTIONS

1. Make the first node (or highest qualifier) of your dataset name the name of your usercatalog. Example: GPVS1.yourfile.cl.
2. Specify names for data and index components.
3. Specify the appropriate suffix to your dataset name by dataset type:

CLUSTER	gpvs1.yourfile.CL
DATA	gpvs1.yourfile.DA
INDEX	gpvs1.yourfile.IX
ALTERNATE INDEX	gpvs1.yourfile.AIX
PATH	gpvs1.yourfile.PATH

DATA CISIZE

If the record length is not greater than 200 bytes, then use 2048 as DATA CISIZE. For record lengths greater than 200 bytes, 4096 will generally be the best DATA CISIZE. In a CICS environment, avoid record length greater than 400 bytes, and never use DATA CISIZES greater than 4096.

INDEX CISIZE

The best strategy here is to obtain a two level index. Generally a 512 INDEX CISIZE will achieve this, especially if the size of the dataset is 50 cylinders or less. For larger datasets, 1024 may better achieve the two-level index.

BUFFERSPACE

For direct or random processing, multiply the DATA CISIZE by 2, then round up to the next multiple of 2048. For sequential processing, add 6 times the DATA CISIZE plus 4 times the INDEX CISIZE.

The following parameters are not "define time" parameters, but are necessary for CICS files being entered in the file control tables, and may be considered "execution time" JCL parameters for optimizing VSAM batch performance.

STRNO

The number of concurrent accesses to the dataset.

BUFNI

The number of index buffers; minimum value is equal to STRNO.

BUFND

The number of data buffers; minimum value is equal to STRNO + 1.

For help with your individual applications, or questions regarding this Bulletin, please call Phil Stephenson at 533-4292, or Mike Berrett at 533-3260.